

# Using Job Arrays to Group Related Jobs

---

A job array is a collection of jobs that differ from each other by only a single index parameter. Creating a job array provides an easy way to group related jobs together. For example, if you have a parameter study that requires you to run your application five times, each with a different input parameter, you can use a job array instead of creating five separate PBS scripts and submitting them separately.

## Creating a Job Array

To create a job array, use a single PBS script and use the -J option to specify a range for the index parameter, either on the qsub command line or within your PBS script. For example, if you submit the following script, a job array with five sub-jobs will be created:

```
#PBS -lselect=4:ncpus=20:model=ivy
#PBS -lwalltime=8:00:00
#PBS -J 1-5

# For each sub-job, a value provided by the -J
# option is used for PBS_ARRAY_INDEX

mkdir dir.$PBS_ARRAY_INDEX
cd dir.$PBS_ARRAY_INDEX

mpiexec ../a.out < ../input.$PBS_ARRAY_INDEX
```

Submitting the script to PBS will return a PBS\_ARRAY\_ID. For example:

```
% qsub job_array_script
28846[].pbspl1.nas.nasa.gov
```

Each sub-job in this job array will have a PBS\_JOBID that includes both the PBS\_ARRAY\_ID and a unique PBS\_ARRAY\_INDEX value within the brackets. For example:

```
28846[1].pbspl1.nas.nasa.gov
28846[2].pbspl1.nas.nasa.gov
28846[3].pbspl1.nas.nasa.gov
28846[4].pbspl1.nas.nasa.gov
28846[5].pbspl1.nas.nasa.gov
```

When you develop your script, note the following:

- Interactive submission of job arrays is not allowed.
- Job arrays are required to be re-runnable. Submitting jobs with the -r n option is not allowed.
- The maximum array size is limited to 5 for array jobs submitted to the devel queue. Only one sub-job of the array job can be running in the devel queue.
- For all other queues, the maximum array size is limited to 501.

## Checking Status

The status of the sub-jobs is not displayed by default. For example, the following qstat options shows the job array as a single job:

```
% qstat -a
or
% qstat -nu your_username
or
% qstat -J
```

JobID	User	Queue	Jobname	TSK	Nds	wallt	S
-----							
28846[].pbspl1	user	normal	myscript	20	4	00:02	B

In the example above, the "B" status applies only to job arrays. This status indicates that at least one sub-job has left the "Q" (queued) state and is running or has run, but not all sub-jobs have run.

To check the status of the sub-jobs, use either the -Jt option or the -t option with an array specified. For example:

```
% qstat -Jt
or
% qstat -t 28846[].pbspl1.nas.nasa.gov
```

JobID	User	Queue	Jobname	TSK	Nds	wallt	S
-----							
28846[1].pbspl1	user	normal	myscript	20	4	00:02	R
28846[2].pbspl1	user	normal	myscript	20	4	00:02	R
28846[3].pbspl1	user	normal	myscript	20	4	00:02	Q
28846[4].pbspl1	user	normal	myscript	20	4	00:02	Q
28846[5].pbspl1	user	normal	myscript	20	4	00:02	Q

## Deleting a Job Array or Sub-Job

To delete a job array or a sub-job, use the `qdel` command and specify the array or sub-job. For example:

```
% qdel "28846[]"
```

```
% qdel "28846[5]"
```

## More Information

For more information about job arrays, see Chapter 8 of the PBS Professional User's Guide, which is available in the `/PBS/doc` directory on Pleiades.

---

Article ID: 464

Last updated: 31 Jan, 2020

Updated by: Moyer M.

Revision: 17

Optimizing/Troubleshooting -> PBS Jobs -> Using Job Arrays to Group Related Jobs

<https://www.nas.nasa.gov/hecc/support/kb/entry/464/>